

L Number	Hits	Search Text	DB	Time stamp
-	1501	707/2.ccls.	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/11/20 13: 25
-	37	707/2.ccls. and (query adj cost)	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/11/19 10: 02
-	34	(707/2.ccls. and (query adj cost)) and (@rlad <=20010228 @ad<=20010228)	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/11/21 12: 48
-	78	707/2.ccls. and (generat\$3 adj index) and (@rlad <=20010228 @ad<=20010228)	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/11/19 10: 09
-	50	(707/2.ccls. and (generat\$3 adj index) and (@rlad <=20010228 @ad<=20010228)) and cost	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/11/19 10: 18
-	1	09/911,784	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/11/20 13: 24
-	1	6356889.URPN.	USPAT	2003/11/19 12: 49
-	9	("5276870" "5822747" "5822750" "5864840" "5864842" "5930785" "6122644" "6199063" "6240406").PN.	USPAT	2003/11/19 12: 52
-	8	09/139843	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/11/19 13: 31
-	6	09/087617	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/11/19 13: 31
-	15	"5950186"	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/11/19 15: 48
-	14	"5960423"	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/11/19 15: 51
-	11	"5913207"	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/11/19 15: 51
-	15	707/2.ccls. and (dynamic adj index\$3)	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/11/20 08: 54
-	35	707/2,3,4,5.ccls. and (dynamic adj index\$3)	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/11/20 08: 51
-	32	(707/2,3,4,5.ccls. and (dynamic adj index\$3)) and (@rlad <=20010228 @ad<=20010228)	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/11/20 08: 54
-	27	707/2.ccls. and (delet\$3 adj index\$3)	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/11/20 08: 54
-	27	(707/2.ccls. and (delet\$3 adj index\$3)) and (@rlad <=20010228 @ad<=20010228)	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/11/20 12: 11
-	54	5404510.URPN.	USPAT	2003/11/20 09: 06
-	80	(index adj maintenance) and (@rlad <=20010228 @ad<=20010228)	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/11/20 15: 33
-	39	((index adj maintenance) and (@rlad <=20010228 @ad<=20010228)) and cost	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/11/20 12: 11



	55	"5404510"	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/11/21 09:32
	14	5778353.URPN.	USPAT	2003/11/21 09:30
	70	partial adj index	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/11/21 09:32
	57	(partial adj index) and (@rlad <=20010228 @ad<=20010228)	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/11/21 10:37
	11	("5778353" "5897632" "5913206" "5913207" "5926813" "5950186" "5960423" "5991754" "6026390" "6134543" "6169983").PN.	USPAT	2003/11/21 10:29
	3	6169983.URPN.	USPAT	2003/11/21 10:29
	7	("5440730" "5644763" "5685003" "5737732" "5835959" "5950186" "6014663").PN.	USPAT	2003/11/21 10:30
	7	6223171.URPN.	USPAT	2003/11/21 10:35
	92	(merg\$3 adj index) and (@rlad <=20010228 @ad<=20010228)	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/11/21 10:37
	0	((merg\$3 adj index) and (@rlad <=20010228 @ad<=20010228)) and 707/1,2,3,4,10	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/11/21 10:38
	45	((merg\$3 adj index) and (@rlad <=20010228 @ad<=20010228)) and 707/1,2,3,4,10,100,101.ccls.	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/11/21 12:35
	1	6275822.URPN.	USPAT	2003/11/21 10:41
	4	("4644471" "5239663" "5619692" "5926807").PN.	USPAT	2003/11/21 10:42
	410	(index adj selection) and (@rlad <=20010228 @ad<=20010228)	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/11/21 12:36
	141	((index adj selection) and (@rlad <=20010228 @ad<=20010228)) and cost	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/11/21 12:36
	49	generat\$3 adj new adj index	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/11/21 12:52
	42	(generat\$3 adj new adj index) and (@rlad <=20010228 @ad<=20010228)	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/11/21 12:53
	103	creat\$3 adj (new second) adj index	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/11/21 12:53
	89	(creat\$3 adj (new second) adj index) and (@rlad <=20010228 @ad<=20010228)	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/11/21 12:53
	22	("5276870" "5301317" "5335345" "5347653" "5367675" "5404510" "5495605" "5555409" "5560007" "5608904" "5630120" "5666528" "5671403" "5675785" "5694591" "5745904" "5758145" "5765147" "5765168" "5943666" "5960428" "6006220").PN.	USPAT	2003/11/21 13:06
	15	5727196.URPN.	USPAT	2003/11/21 13:54
	3	("4468732" "4956774" "5043872").PN.	USPAT	2003/11/21 14:03

## On the Selection of Secondary Indices in Relational Databases (1993) (Make Corrections) (6 citations)

Sunil Choenni, Henk M. Blanken, Thiel Chang  
Data Knowledge Engineering

**CiteSeer** [Home/Search](#) [Bookmark](#) [Context](#)

[Related](#)

[\(Enter summary\)](#)

View or download:

[wwwis.cs.utwente.n...e93.accepted.ps.gz](#)

Cached: [PS.gz](#) [PS](#) [PDF](#) [DjVu](#) [Image](#) [Update](#) [Help](#)

From: [wwwis.cs.utwente.n...journalpaper](#) (more)

Homepages: [S.Choenni](#) [H.Blanken](#) [2]

[T.Chang](#) [2] [3] [4] [HPSearch](#) (Update Links)

Rate this article: [1](#) [2](#) [3](#) [4](#) [5](#) (best)

[Comment on this article](#)

**Abstract:** An important problem in the physical design of databases is the selection of secondary indices. In general, this problem can not be solved in an optimal way due to the complexity of the selection process. Often use is made of heuristics such as the well-known ADD and DROP algorithms. In this paper it will be shown that frequently used cost functions can be classified as super- or submodular functions. For these functions several mathematical properties have been derived which reduce the... [\(Update\)](#)

Context of citations to this paper: [More](#)

...as adding a set of features to the database. In particular, the index selection part of our VIS problem has been well studied [FST88,CBC93] in the context of physical database design. Choosing indexes for materialized views is a straightforward extension. What is...

.... see among others [BaPS 90, Card 75, FiST 88, IpSR 83, Schk 75, Whan 87] and still research is going on this topic [ChBC 93a, ChBC 93b, Roze 93] The idea of our approach, in the context of object oriented databases, is to determine a number of non overlapping paths....

Cited by: [More](#)

Facilitating Hard Active Database Applications - Warshaw (2001) [\(Correct\)](#)

Datplex: An Extensible Design-Aid Tool For Physical Database.. - Lu, Ooi, Qiu, Tan [\(Correct\)](#)

Physical Database Design for Data Warehouses - Labio, Quass, Adelberg (1997) [\(Correct\)](#)

Active bibliography (related documents): [More](#) [All](#)

1.2: Index Configurations in Object-Oriented Databases - Choenni, Bertino, Blanken.. [\(Correct\)](#)

1.0: Index Selection in Relational Databases - Choenni, Blanken, Chang (1993) [\(Correct\)](#)

0.7: A Genetic Algorithm for the Index Selection - Problem Jozef Kratica [\(Correct\)](#)

Similar documents based on text: [More](#) [All](#)

0.4: Predicting the cost-quality trade-off for.. - Blok, Hiemstra.. [\(Correct\)](#)

0.1: A Selectivity Model for Fragmented Relations in.. - Blok, Choenni.. (2001) [\(Correct\)](#)

0.1: On the Selection of Optimal Index Configuration in OO.. - Choenni, Bertino.. (1994) [\(Correct\)](#)

Related documents from co-citation: [More](#) [All](#)

5: Physical Database Design for Relational Databases (context) - Finkelstein, Schkolnick et al. - 1988

4: A framework for automating physical databasedesign (context) - Rozen, Shasha - 1991

3: Optimal update policies for distribtued materialized views (context) - Segev, Fang - 1991

BibTeX entry: [\(Update\)](#)

S. Choenni, H. Blanken, and T. Chang. On the selection of secondary indices in relational databases. Data and Knowledge Engineering, 11:207-33, 1993. <http://citeseer.nj.nec.com/choenni93selection.html> [More](#)

```
@article{ choenni93selection,
  author = "Sunil Choenni and Henk M. Blanken and T. Chang",
  title = "On the Selection of Secondary Indices in Relational Databases",
  journal = "Data Knowledge Engineering",
  volume = "11",
  number = "3",
  pages = "207 -",
  year = "1993",
```

Searching for **PHRASE index selection**.

Restrict to: [Header](#) [Title](#) Order by: [Citations](#) [Hubs](#) [Usage](#) [Date](#) Try: [Amazon](#) [B&N](#) [Google \(RI\)](#) [Google \(Web\)](#) [CSB](#) [DBLP](#)

99 documents found. Order: **citations weighted by year**.

[Index Selection for OLAP - Gupta, Harinarayan, Rajaraman, Ullman \(1997\) \(Correct\) \(68 citations\)](#)

**Index Selection** for OLAP #Himanshu Gupta Venky Harinarayan

This paper is the #rst to explore the **index-selection** problem and automate it with provably

Now consider integrating the subcube selection and **index selection** into one step. We use the

[www.cise.ufl.edu/~jgreenbe/research/..papers/23.pdf](http://www.cise.ufl.edu/~jgreenbe/research/..papers/23.pdf)

[Cluster-based Language Models for Distributed Retrieval - Xu, Croft \(1999\) \(Correct\) \(29 citations\)](#)

data structure is called the collection **selection index** #23#Collection selection consists of in the same time period. The collection **selection index** summarizes each collection as a whole. In

is illustrated by Figure 1. Collection **Selection Index** collection 1 collection 2 collection3

[www.cs.virginia.edu/~cyberia/cs851/papers/xu-clstr.pdf](http://www.cs.virginia.edu/~cyberia/cs851/papers/xu-clstr.pdf)

[Cluster-based Language Models For Distributed Retrieval - Xu, Croft \(1999\) \(Correct\) \(25 citations\)](#)

data structure is called the collection **selection index** [23]Collection selection consists of

in the same time period. The collection **selection index** summarizes each collection as a whole. In

is illustrated by Figure 1. Collection **Selection Index** collection 1 collection 2 collection3

[ciir.cs.umass.edu/pubfiles/ir-153.ps](http://ciir.cs.umass.edu/pubfiles/ir-153.ps)

[Comparing the Performance of Database Selection.. - French, Powell, Callan, .. \(1999\) \(Correct\) \(15 citations\)](#)

search, the CORI approach creates a database **selection index** in which each database is represented by its

[www.cs.virginia.edu/~cyberia/papers/SIGIR99.ps](http://www.cs.virginia.edu/~cyberia/papers/SIGIR99.ps)

[Subtopic Structuring for Full-Length Document Access - Hearst, Plaunt \(1993\) \(Correct\) \(55 citations\)](#)

can im- prove on categorization and **index selection** tasks. We then describe an experiment in

[www.ai.mit.edu/people/jimmylin/papers/Hearst93.pdf](http://www.ai.mit.edu/people/jimmylin/papers/Hearst93.pdf)

[Optimization of Run-time Management of Data Intensive Web .. - Florescu, Levy, Suciu.. \(1999\) \(Correct\) \(12 citations\)](#)

as view materialization [24, 13, 12, 14, 6]**index selection**, data caching [16, 15, 8]multiple query

14, 6]multiple query optimization [23] and **index selection**. All these techniques are of course

[www.research.att.com/~suciu/strudel/external/files/\\_F269651526.ps](http://www.research.att.com/~suciu/strudel/external/files/_F269651526.ps)

[Optimization of Run-time Management of Data Intensive.. - Florescu, Levy, Suciu.. \(1999\) \(Correct\) \(12 citations\)](#)

view materialization [13, 23, 10, 12, 11, 6]**index selection** [7]function caching [16, 14]multiple

12, 11, 6]multiple query optimization [22]**index selection** [7]and parameterized query optimization

[rodin.inria.fr/dataFiles/FLSY99b.ps](http://rodin.inria.fr/dataFiles/FLSY99b.ps)

[Effective Retrieval with Distributed Collections - Xu, Callan \(1998\) \(Correct\) \(19 citations\)](#)

is to use phrase information in the collection **selection index** and the other is query expansion. Both

8/98 \$5.00. creates a collection **selection index**. The collection **selection index** consists of

a collection **selection index**. The collection **selection index** consists of a set of virtual documents,

[www.cs.umass.edu/~xu/sigir98-final.ps](http://www.cs.umass.edu/~xu/sigir98-final.ps)

[Physical Database Design for Data Warehouses - Labio, Quass, Adelberg \(1997\) \(Correct\) \(17 citations\)](#)

the total down time. We call this the view **index selection** (VIS) problem. We present an exhaustive

materialized views, view maintenance, **index selection**, and physical database design. 1

views it is necessary to consider the view selection and **index selection** together. If view selection is

[ftp.dblab.ntua.gr/pub/dwq/view.ps.gz](http://ftp.dblab.ntua.gr/pub/dwq/view.ps.gz)

[The Dimensional Fact Model: A Conceptual Model For Data.. - Golfarelli, Maio, Rizzi \(1998\) \(Correct\) \(10 citations\)](#)

such as materialization of views 2,15 and **index selection** 13,16 no significant effort has been

13. H. Gupta, V. Harinarayan and A. Rajaraman, **Index selection** for OLAP, Proc. Int. Conf. Data Engineering,

ftp-db.deis.unibo.it/pub/stefano/ijcis98.ps

A Comprehensive Approach to Horizontal Class Fragmentation in ... - Ezeife, Barker (1995) (Correct) (18 citations)  
of data. The work of Bertino and Kim [1] on **index selection** is also complementary to our work, but  
[www.cs.uwindsor.ca/users/c/cezeife/journal1.ps](http://www.cs.uwindsor.ca/users/c/cezeife/journal1.ps)

Rethinking Database System Architecture: Towards a.. - Chaudhuri, Weikum (2000) (Correct) (3 citations)  
of the application. These knobs include **index selection**, data placement across parallel disks, and  
[www-dbs.cs.uni-sb.de/public\\_html/papers/cairo-final.pdf](http://www-dbs.cs.uni-sb.de/public_html/papers/cairo-final.pdf)

A Methodological Framework for Data Warehouse Design - Golfarelli, Rizzi (1998) (Correct) (7 citations)  
[1] 8]materialization of views [2] 9] and **index selection** [7] 10]no significant effort has been  
provided by the DBMS to be taken into account. **Index selection** has a crucial role in determining the DW  
[ftp-db.deis.unibo.it/pub/stefano/dolap98.ps.gz](http://ftp-db.deis.unibo.it/pub/stefano/dolap98.ps.gz)

First 20 documents Next 20

Try your query at: [Amazon](#) [Barnes & Noble](#) [Google \(RI\)](#) [Google \(Web\)](#) [CSB](#) [DBLP](#)

CiteSeer - [citeseer.org](#) - [Terms of Service](#) - [Privacy Policy](#) - Copyright © 1997-2002 NEC Research Institute



[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

Search:  The ACM Digital Library  The Guide

"dynamic index"

**SEARCH**

THE ACM DIGITAL LIBRARY

[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used [dynamic index](#)

Found 306 of 122,783

Sort results by

relevance

[Save results to a Binder](#)

Try an [Advanced Search](#)

Display results

expanded form

[Search Tips](#)  
 Open results in a new window

Try this search in [The ACM Guide](#)

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale

### 1 [An efficient expected cost algorithm for dynamic indexing of spatial objects](#)

Gopi K. Attaluri

October 1994 **Proceedings of the 1994 conference of the Centre for Advanced Studies on Collaborative research**

Full text available: [pdf\(173.67 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)



Spatial object management is important in numerous application areas, including geographic data, CAD, and VLSI design. It often involves insertion, deletion, and range search of spatial objects, so requires efficient dynamic indexing of objects. This paper describes an (in memory) indexing algorithm developed in the context of concurrency control for large, unstructured objects in non-traditional database applications. It supports insertion, deletion, and range search of orthogonal spatial object ...

**Keywords:** concurrency control, indexing, range locking, runtime complexity, spatial databases

### 2 [XML indexing and compression: ViST: a dynamic index method for querying XML data by tree structures](#)

Haixun Wang, Sanghyun Park, Wei Fan, Philip S. Yu

June 2003 **Proceedings of the 2003 ACM SIGMOD international conference on Management of data**

Full text available: [pdf\(244.47 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)



With the growing importance of XML in data exchange, much research has been done in providing flexible query facilities to extract data from structured XML documents. In this paper, we propose ViST, a novel index structure for searching XML documents. By representing both XML documents and XML queries in structure-encoded sequences, we show that querying XML data is equivalent to finding subsequence matches. Unlike index methods that disassemble a query into multiple sub-queries, and then *join* ...

### 3 [Segment indexes: dynamic indexing techniques for multi-dimensional interval data](#)

Curtis P. Kolovson, Michael Stonebraker

April 1991 **ACM SIGMOD Record, Proceedings of the 1991 ACM SIGMOD international conference on Management of data**, Volume 20 Issue 2

Full text available: [pdf\(1.02 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)



4 Physical storage structures: The K-D-B-tree: a search structure for large multidimensional dynamic indexes

John T. Robinson

April 1981 **Proceedings of the 1981 ACM SIGMOD international conference on Management of data**

Full text available: [pdf\(723.91 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

The problem of retrieving multikey records via range queries from a large, dynamic index is considered. By *large* it is meant that most of the index must be stored on secondary memory. By *dynamic* it is meant that insertions and deletions are intermixed with queries, so that the index cannot be built beforehand. A new data structure, the *K-D-B-tree*, is presented as a solution to this problem. K-D-B-trees combine properties of K-D-trees and B-trees. It is expected that the mult ...

5 Physical database design: R-trees: a dynamic index structure for spatial searching

Antonin Guttman

June 1984 **Proceedings of the 1984 ACM SIGMOD international conference on Management of data**

Full text available: [pdf\(1.07 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

In order to handle spatial data efficiently, as required in computer aided design and geo-data applications, a database system needs an index mechanism that will help it retrieve data items quickly according to their spatial locations. However, traditional indexing methods are not well suited to data objects of non-zero size located in multi-dimensional spaces. In this paper we describe a dynamic index structure called an R-tree which meets this need, and give algorithms for searching and updating ...

6 Comparison of access methods for time-evolving data

Betty Salzberg, Vassilis J. Tsotras

June 1999 **ACM Computing Surveys (CSUR)**, Volume 31 Issue 2

Full text available: [pdf\(529.53 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper compares different indexing techniques proposed for supporting efficient access to temporal data. The comparison is based on a collection of important performance criteria, including the space consumed, update processing, and query time for representative queries. The comparison is based on worst-case analysis, hence no assumptions on data distribution or query frequencies are made. When a number of methods have the same asymptotic worst-case behavior, features in the methods tha ...

**Keywords:** I/O performance, access methods, structures, temporal databases

7 Multidimensional access methods

Volker Gaede, Oliver Günther

June 1998 **ACM Computing Surveys (CSUR)**, Volume 30 Issue 2

Full text available: [pdf\(1.05 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Search operations in databases require special support at the physical level. This is true for conventional databases as well as spatial databases, where typical search operations include the point query (find all objects that contain a given search point) and the region query (find all objects that overlap a given search region). More than ten years of spatial database research have resulted in a great variety of multidimensional access methods to support ...

**Keywords:** data structures, multidimensional access methods

**8 The SR-tree: an index structure for high-dimensional nearest neighbor queries**

Norio Katayama, Shin'ichi Satoh

June 1997 **ACM SIGMOD Record , Proceedings of the 1997 ACM SIGMOD international conference on Management of data**, Volume 26 Issue 2Full text available: [pdf\(1.41 MB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Recently, similarity queries on feature vectors have been widely used to perform content-based retrieval of images. To apply this technique to large databases, it is required to develop multidimensional index structures supporting nearest neighbor queries efficiently. The SS-tree had been proposed for this purpose and is known to outperform other index structures such as the R\*-tree and the K-D-B-tree. One of its most important features is that it employs bounding spheres rather than boundi ...

**9 On packing R-trees**

Ibrahim Kamel, Christos Faloutsos

December 1993 **Proceedings of the second international conference on Information and knowledge management**Full text available: [pdf\(929.84 KB\)](#)Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)**10 Parallel R-trees**

Ibrahim Kamel, Christos Faloutsos

June 1992 **ACM SIGMOD Record , Proceedings of the 1992 ACM SIGMOD international conference on Management of data**, Volume 21 Issue 2Full text available: [pdf\(991.13 KB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We consider the problem of exploiting parallelism to accelerate the performance of spacial access methods and specifically, R-trees [11]. Our goal is to design a server for spatial data, so that to maximize the throughput of range queries. This can be achieved by (a) maximizing parallelism for large range queries, and (b) by engaging as few disks as possible on point queries [22]. We propose a simple hardware architecture consisting of one processor with several disks attached to ...

**11 High performance clustering based on the similarity join**

Christian Böhm, Bernhard Braunschmidt, Markus Breunig, Hans-Peter Kriegel

November 2000 **Proceedings of the ninth international conference on Information and knowledge management**Full text available: [pdf\(134.57 KB\)](#)Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

**Keywords:** clustering, data mining, database primitives, multidimensional index structure, similarity join

**12 Searching spatial objects with index by dimensional projections**

Xiaoming Cheng, Huizhu Lu, G. E. Hedrick

April 1992 **Proceedings of the 1992 ACM/SIGAPP Symposium on Applied computing: technological challenges of the 1990's**Full text available: [pdf\(588.72 KB\)](#)Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

**13 Special issue on spatial database systems: An introduction to spatial database systems**

Ralf Hartmut Güting

October 1994 **The VLDB Journal — The International Journal on Very Large Data Bases**,

Volume 3 Issue 4

Full text available: [pdf\(2.50 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

We propose a definition of a spatial database system as a database system that offers spatial data types in its data model and query language, and supports spatial data types in its implementation, providing at least spatial indexing and spatial join methods. Spatial database systems offer the underlying database technology for geographic information systems and other applications. We survey data modeling, querying, data structures and algorithms, and system architecture for such systems. The em ...

**14 Redundancy in spatial databases**

J. A. Orenstein

June 1989 **ACM SIGMOD Record , Proceedings of the 1989 ACM SIGMOD international conference on Management of data**, Volume 18 Issue 2

Full text available: [pdf\(1.37 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Spatial objects other than points and boxes can be stored in spatial indexes, but the techniques usually require the use of approximations that can be arbitrarily bad. This leads to poor performance and highly inaccurate responses to spatial queries. The situation can be improved by storing some objects in the index redundantly. Most spatial indexes permit no flexibility in adjusting the amount of redundancy. Spatial indexes based on z-order permit this flexibility. Accuracy of the query re ...

**15 Scalable integrated region-based image retrieval using IRM and statistical clustering**

James Z. Wang, Yanping Du

January 2001 **Proceedings of the first ACM/IEEE-CS joint conference on Digital libraries**

Full text available: [pdf\(1.73 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Statistical clustering is critical in designing scalable image retrieval systems. In this paper, we present a scalable algorithm for indexing and retrieving images based on region segmentation. The method uses statistical clustering on region features and IRM (Integrated Region Matching), a measure developed to evaluate overall similarity between images that incorporates properties of all the regions in the images by a region-matching scheme. Compared with retrieval based on individual ...

**Keywords:** clustering, content-based image retrieval, integrated region matching, segmentaton, wavelets

**16 On effective multi-dimensional indexing for strings**

H. V. Jagadish, Nick Koudas, Divesh Srivastava

May 2000 **ACM SIGMOD Record , Proceedings of the 2000 ACM SIGMOD international conference on Management of data**, Volume 29 Issue 2

Full text available: [pdf\(1.15 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

As databases have expanded in scope from storing purely business data to include XML documents, product catalogs, e-mail messages, and directory data, it has become increasingly important to search databases based on wild-card string matching: prefix matching, for example, is more common (and useful) than exact matching, for such data. In many cases, matches need to be on multiple attributes/dimensions, with correlations

between the dimensions. Traditional multi-dimensional index structures, ...

**17 Simple QSF-trees: an efficient and scalable spatial access method**

Byunggu Yu, Ratko Orlandic, Martha Evens

November 1999 **Proceedings of the eighth international conference on Information and knowledge management**

Full text available:  [pdf\(1.33 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The development of high-performance spatial access methods that can support complex operations of large spatial databases continues to attract considerable attention. This paper introduces QSF-trees, an efficient and scalable structure for indexing spatial objects, which has some important advantages over R\*-trees. QSF-trees eliminate overlapping of index regions without forcing object clipping or sacrificing the select ...

**Keywords:** database management, point access methods, spatial access methods, spatial database, topological relations

**18 Distance browsing in spatial databases**

Gísli R. Hjaltason, Hanan Samet

June 1999 **ACM Transactions on Database Systems (TODS)**, Volume 24 Issue 2

Full text available:  [pdf\(460.81 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We compare two different techniques for browsing through a collection of spatial objects stored in an R-tree spatial data structure on the basis of their distances from an arbitrary spatial query object. The conventional approach is one that makes use of a k-nearest neighbor algorithm where k is known prior to the invocation of the algorithm. Thus if  $m < k$  neighbors are needed, the k-nearest neighbor alg ...

**Keywords:** R-trees, distance browsing, hierarchical spatial data structures, nearest neighbors, ranking

**19 Indexing large metric spaces for similarity search queries**

Tolga Bozkaya, Meral Ozsoyoglu

September 1999 **ACM Transactions on Database Systems (TODS)**, Volume 24 Issue 3

Full text available:  [pdf\(281.78 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

One of the common queries in many database applications is finding approximate matches to a given query item from a collection of data items. For example, given an image database, one may want to retrieve all images that are similar to a given query image. Distance-based index structures are proposed for applications where the distance computations between objects of the data domain are expensive (such as high-dimensional data) and the distance function is metric. In this paper we consider ...

**20 Indexing schemes for random points**

Elias Koutsoupias, David Taylor

January 1999 **Proceedings of the tenth annual ACM-SIAM symposium on Discrete algorithms**

Full text available:  [pdf\(838.33 KB\)](#)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

ACM Portal is published by the Association for Computing Machinery. Copyright © 2003 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)



[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

Search:  The ACM Digital Library  The Guide

query <and> cost <and> index

**SEARCH**

THE ACM DIGITAL LIBRARY



[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used query and cost and index

Found 122,783 of 122,783

Sort results by

relevance

Save results to a Binder

Display results

expanded form

Search Tips  
 Open results in a new window

Try an [Advanced Search](#)

Try this search in [The ACM Guide](#)

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale

### 1 [A cost model for query processing in high dimensional data spaces](#)

Christian Böhm

June 2000 **ACM Transactions on Database Systems (TODS)**, Volume 25 Issue 2

Full text available: [pdf\(362.22 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)



During the last decade, multimedia databases have become increasingly important in many application areas such as medicine, CAD, geography, and molecular biology. An important research topic in multimedia databases is similarity search in large data sets. Most current approaches that address similarity search use the feature approach, which transforms important properties of the stored objects into points of a high-dimensional space (feature vectors). Thus, similarity search is transformed ...

**Keywords:** cost model, multidimensional index

### 2 [Minimum cost selection of secondary indexes for formatted files](#)

Henry D. Anderson, P. Bruce Berra

March 1977 **ACM Transactions on Database Systems (TODS)**, Volume 2 Issue 1

Full text available: [pdf\(1.74 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)



Secondary indexes are often used in database management systems for secondary key retrieval. Although their use can improve retrieval time significantly, the cost of index maintenance and storage increases the overhead of the file processing application. The optimal set of indexed secondary keys for a particular application depends on a number of application dependent factors. In this paper a cost function is developed for the evaluation of candidate indexing choices and applied to the opti ...

**Keywords:** Boolean query, access methods, access path, cost function, data management, database, file design, file organization, inverted file, inverted index, maintenance, optimization, retrieval, secondary index, secondary key, secondary key access

### 3 [Query processing: Factorizing complex predicates in queries to exploit indexes](#)

Surajit Chaudhuri, Prasanna Ganesan, Sunita Sarawagi

June 2003 **Proceedings of the 2003 ACM SIGMOD international conference on Management of data**



Full text available: [pdf\(240.56 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Decision-support applications generate queries with complex predicates. We show how the *factorization* of complex query expressions exposes significant opportunities for exploiting available indexes. We also present a novel idea of relaxing predicates in a complex condition to create possibilities for factoring. Our algorithms are designed for easy integration with existing query optimizers and support multiple optimization levels, providing different trade-offs between plan complexity and ...

4 Query processing: A characterization of the sensitivity of query optimization to storage access cost parameters 

Frederick R. Reiss, Tapas Kanungo

June 2003 **Proceedings of the 2003 ACM SIGMOD international conference on on Management of data**

Full text available: [pdf\(255.35 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Most relational query optimizers make use of information about the costs of accessing tuples and data structures on various storage devices. This information can at times be off by several orders of magnitude due to human error in configuration setup, sudden changes in load, or hardware failure. In this paper, we attempt to answer the following questions:• Are inaccurate access cost estimates likely to cause a typical query optimizer to choose a suboptimal query plan?• If an optimizer ...

**Keywords:** autonomic computing, computational geometry, databases, parametric query optimization, storage systems

5 Location awareness and moving objects: Efficient placement of geographical data over broadcast channel for spatial range query under quadratic cost model 

Jianting Zhang, Le Gruenwald

September 2003 **Proceedings of the 3rd ACM international workshop on Data engineering for wireless and mobile access**

Full text available: [pdf\(326.37 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Data broadcasting is well known for its excellent scalability. Most geographical data, such as weather and traffic, is public information that has a large amount of potential users which makes it very suitable for broadcast. The query response time is greatly affected by the order in which data items are being broadcast. This paper proposes an efficient method to place geographical data items over broadcast channel that reduces access time for spatial range queries on them. This paper then perfo ...

**Keywords:** cost model, data broadcast, geographical information, mobile computing, optimization, query processing

6 A cost model for similarity queries in metric spaces 

Paolo Ciaccia, Marco Patella, Pavel Zezula

May 1998 **Proceedings of the seventeenth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems**

Full text available: [pdf\(1.12 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

7 Improved query performance with variant indexes 

Patrick O'Neil, Dallan Quass

June 1997 **ACM SIGMOD Record , Proceedings of the 1997 ACM SIGMOD international conference on Management of data**, Volume 26 Issue 2

Full text available: [pdf\(1.54 MB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The read-mostly environment of data warehousing makes it possible to use more complex indexes to speed up queries than in situations where concurrent updates are present. The current paper presents a short review of current indexing technology, including row-set representation by Bitmaps, and then introduces two approaches we call Bit-Sliced indexing and Projection indexing. A Projection index materializes all values of a column in RID order, and a Bit-Sliced index essentially takes an orth ...

## 8 Indexing large metric spaces for similarity search queries

Tolga Bozkaya, Meral Ozsoyoglu

September 1999 **ACM Transactions on Database Systems (TODS)**, Volume 24 Issue 3Full text available: [pdf\(281.78 KB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

One of the common queries in many database applications is finding approximate matches to a given query item from a collection of data items. For example, given an image database, one may want to retrieve all images that are similar to a given query image. Distance-based index structures are proposed for applications where the distance computations between objects of the data domain are expensive (such as high-dimensional data) and the distance function is metric. In this paper we consider ...

## 9 Efficiency: Compression of inverted indexes For fast query evaluation

Falk Scholer, Hugh E. Williams, John Yiannis, Justin Zobel

August 2002 **Proceedings of the 25th annual international ACM SIGIR conference on Research and development in information retrieval**Full text available: [pdf\(174.13 KB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Compression reduces both the size of indexes and the time needed to evaluate queries. In this paper, we revisit the compression of inverted lists of document postings that store the position and frequency of indexed terms, considering two approaches to improving retrieval efficiency: better implementation and better choice of integer compression schemes. First, we propose several simple optimisations to well-known integer compression schemes, and show experimentally that these lead to significant ...

**Keywords:** index compression, integer coding, inverted indexes, retrieval efficiency

## 10 Document query processing strategies: cost evaluation and heuristics

E. Bertino, S. Gibbs, F. Rabitti

April 1988 **Conference Sponsored by ACM SIGOIS and IEEECS TC-OA on Office information systems**Full text available: [pdf\(969.16 KB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This paper describes query processing strategies used in a system where queries are specified on formatted data and text components of documents. The system provides different access methods: indexes for formatted data and signature files for the text. Four basic document queries, representative of a wide range of queries, are examined in detail. Strategies for query processing are presented together with an evaluation of the costs of the various strategies. Finally, general heuristics ...

## 11 The onion technique: indexing for linear optimization queries

Yuan-Chi Chang, Lawrence Bergman, Vittorio Castelli, Chung-Sheng Li, Ming-Ling Lo, John R. Smith

May 2000 **ACM SIGMOD Record , Proceedings of the 2000 ACM SIGMOD international**

**conference on Management of data**, Volume 29 Issue 2Full text available:  [pdf\(326.62 KB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

*This paper describes the Onion technique, a special indexing structure for linear optimization queries. Linear optimization queries ask for top- $N$  records subject to the maximization or minimization of linearly weighted sum of record attribute values. Such query appears in many applications employing linear models and is an effective way to summarize representative cases, such as the top-50 ranked colleges. The Onion indexing is based on a geometric property of convex hull, which guarantees ...*

**Keywords:** *database indexing, linear optimization*

**12 Optimization of object-oriented recursive queries using cost-controlled strategies** 

Rosana S. G. Lanzelotte, Patrick Valduriez, Mohamed Zaït

June 1992 **ACM SIGMOD Record, Proceedings of the 1992 ACM SIGMOD international conference on Management of data**, Volume 21 Issue 2Full text available:  [pdf\(1.07 MB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Object-oriented data models are being extended with recursion to gain expressive power. This complicates the optimization problem which has to deal with recursive queries on complex objects. Because unary operations invoking methods or path expressions on objects may be costly to execute, traditional heuristics for optimizing recursive queries are no longer valid. In this paper we propose a cost-based optimization method which handles object-oriented recursive queries. In particular, it is ...

**13 Query execution in prism and seaview: a cost analysis** 

Brajendra Panda, William Perrizo

February 1995 **Proceedings of the 1995 ACM symposium on Applied computing**Full text available:  [pdf\(699.80 KB\)](#)Additional Information: [full citation](#), [references](#), [index terms](#)

**Keywords:** *data and user classification, military security policies, multilevel databases, query processing*

**14 Web services and performance evaluation: Indexing web access-logs for pattern queries** 

Alexandros Nanopoulos, Yannis Manolopoulos, Maciej Zakrzewicz, Tadeusz Morzy

November 2002 **Proceedings of the fourth international workshop on Web information and data management**Full text available:  [pdf\(187.24 KB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

In this paper, we develop a new indexing method for large web access-logs. We are concerned with pattern queries, which advocate the search for access sequences that contain certain query patterns. This kind of queries find applications in processing web-log mining results (e.g., finding typical/atypical access-sequences). The proposed method focuses on scalability to web-logs' sizes. For this reason, we examine the gains due to signature-trees, which can further improve the scalability to very ...

**15 Cost-based query optimization for metadata repositories** 

Martin Staudt, René Soiron, Christoph Quix, Matthias Jarke

March 1999 **ACM SIGAPP Applied Computing Review**, Volume 7 Issue 2

Full text available: [pdf\(1.04 MB\)](#) Additional Information: [full citation](#), [abstract](#), [index terms](#)

Query optimization strategies for repository systems must take into account the rich and often unpredictable structure of metadata, as well as supporting complex analysis of relationships between those structures. This paper describes rationale, design, and system integration of a cost-based query optimizer offered in ConceptBase, a metadata manager that supports these capabilities by a deductive object-oriented data model. In contrast to most implemented DBMS, the optimizer is not based on the ...

**16 Database principles: A mapping mechanism to support bitmap index and other auxiliary structures on tables stored as primary B<sup>+</sup>-trees**

Eugene Inseok Chong, Chuck Freiwald, Anh-Tuan Tran, Jagannathan Srinivasan, Aravind Yalamanchi, Ramkumar Krishnan, Souripriya Das, Mahesh Jagannath, Richard Jiang  
June 2003 **ACM SIGMOD Record**, Volume 32 Issue 2

Full text available: [pdf\(198.55 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#)

Any auxiliary structure, such as a bitmap or a B<sup>+</sup>-tree index, that refers to rows of a table stored as a primary B<sup>+</sup>-tree (e.g., *tables with clustered index* in Microsoft SQL Server, or *index-organized tables* in Oracle) by their physical addresses would require updates due to inherent volatility of those addresses. To address this problem, we propose a mapping mechanism that 1) introduces a single mapping table, with each row holding one key value from the prima ...

**17 Query optimization in a memory-resident domain relational calculus database system**

Kyu-Young Whang, Ravi Krishnamurthy  
March 1990 **ACM Transactions on Database Systems (TODS)**, Volume 15 Issue 1

Full text available: [pdf\(2.46 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We present techniques for optimizing queries in memory-resident database systems. Optimization techniques in memory-resident database systems differ significantly from those in conventional disk-resident database systems. In this paper we address the following aspects of query optimization in such systems and present specific solutions for them: (1) a new approach to developing a CPU-intensive cost model; (2) new optimization strategies for main-memory query processing; (3) new insight into ...

**18 Optimal indexing using near-minimal space**

C. Heeren, H. V. Jagadish, L. Pitt  
June 2003 **Proceedings of the twenty-second ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems**

Full text available: [pdf\(208.43 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We consider the index selection problem. Given either a fixed query workload or an unknown probability distribution on possible future queries, and a bound  $B$  on how much space is available to build indices, we seek to build a collection of indices for which the average query response time is minimized. We give strong negative and positive performance bounds. Let  $m$  be the number of queries in the workload. We show how to obtain with high probability a collection of indices using space ...

**19 Parametric query optimization**

Yannis E. Ioannidis, Raymond T. Ng, Kyuseok Shim, Timos K. Sellis  
May 1997 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 6 Issue 2

Full text available: [pdf\(378.53 KB\)](#) Additional Information: [full citation](#), [abstract](#), [index terms](#)

In most database systems, the values of many important run-time parameters of the system, the data, or the query are unknown at query optimization time. Parametric query

optimization attempts to identify at compile time several execution plans, each one of which is optimal for a subset of all possible values of the run-time parameters. The goal is that at run time, when the actual parameter values are known, the appropriate plan should be identifiable with essentially no overhead. We present a g ...

## 20 Query processing in a multimedia document system

Elisa Bertino, Fausto Rabbiti, Simon Gibbs

January 1988 **ACM Transactions on Information Systems (TOIS)**, Volume 6 Issue 1

Full text available:  [pdf\(2.94 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Query processing in a multimedia document system is described. Multimedia documents are information objects containing formatted data, text, image, graphics, and voice. The query language is based on a conceptual document model that allows the users to formulate queries on both document content and structure. The architecture of the system is outlined, with focus on the storage organization in which both optical and magnetic devices can coexist. Query processing and the different strategies ...

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2003 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)

**This Page is Inserted by IFW Indexing and Scanning  
Operations and is not part of the Official Record**

**BEST AVAILABLE IMAGES**

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- BLACK BORDERS**
- IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- FADED TEXT OR DRAWING**
- BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- SKEWED/SLANTED IMAGES**
- COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- GRAY SCALE DOCUMENTS**
- LINES OR MARKS ON ORIGINAL DOCUMENT**
- REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- OTHER:** \_\_\_\_\_

**IMAGES ARE BEST AVAILABLE COPY.**

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.